Cosine Similarity Search with Multi-Index Hashing

Sepehr Eghbali and Ladan Tahvildari

arXiv:1610.00574v1 [cs.DB] 14 Sep 2016

Abstract—Due to rapid development of the Internet, recent years have witnessed an explosion in the rate of data generation. Dealing with data at current scales brings up unprecedented challenges. From the algorithmic view point, executing existing linear algorithms in information retrieval and machine learning on such tremendous amounts of data incur intolerable computational and storage costs. To address this issue, there is a growing interest to map data points in large-scale datasets to binary codes. This can significantly reduce the storage complexity of large-scale datasets. However, one of the most compelling reasons for using binary codes or any discrete representation is that they can be used as direct indices into a hash table. Incorporating hash table offers fast query execution; one can look up the nearby buckets in a hash table populated with binary codes to retrieve similar items. Nonetheless, if binary codes are compared in terms of the cosine similarity rather than the Hamming distance, there is no fast exact sequential procedure to find the K closest items to the query other than the exhaustive search. Given a large dataset of binary codes and a binary query, the problem that we address is to efficiently find K closest codes in the dataset that yield the largest cosine similarities to the query. To handle this issue, we first elaborate on the relation between the Hamming distance and the cosine similarity. This allows finding the sequence of buckets to check in the hash table. Having this sequence, we propose a multi-index hashing approach that can increase the search speed up to orders of magnitude in comparison to the exhaustive search and even approximation methods such as LSH. We empirically evaluate the performance of the proposed algorithm on real world datasets.

Index Terms—Nearest neighbour search, binary codes, large-scale retrieval, cosine similarity

1 Introduction

Recent years have witnessed an explosion in the quantity and dimensionality of data. To capture this trend, the term *Big data* was coined by researchers to emphasize the scale of current datasets. This trend is reflected in the growing number of available large-scale datasets on the Internet which have become benchmarks for a broad array of fields. For example, in text mining, Google provides N-gram data obtained from over 5 million books [31]. Resources such as ImageNet [10], Tiny Images [46] and Texmex [23], [39] with 14 millions, 80 millions and 1 billion images, respectively, are typical benchmarks for many machine vision tasks. Dealing with such large-scale datasets brings up unprecedented challenges, in turn necessitating new software tools, algorithms and databases for data acquisition, storage, transmission and retrieval [22].

In this paper, we address the problem of query execution for large-scale datasets. We mainly focus on efficiently solving the KNN query which is an important problem with wide range of applications in many fields such as data mining, information retrieval [15], [23] and handling multimedia as well as spacial databases [19], [26]. Given a collection of items, a query item and a measure of similarity, the exact solution of the KNN problem is the set of the K closest items within the dataset to the query. From an algorithmic point of view, solving KNN for a query point $\mathbf{q} \in \mathbb{R}^d$ requires scanning all n items in the dataset $\mathcal{X} = \{\mathbf{x}_i \in \mathbb{R}^d\}_{i=1}^n$. Assuming that computing similarity of two points in $\mathbb{R}^d$ can be performed in $O(d)$, then KNN can be solved in $O(dn)$ using exhaustive search. Often, executing linear algorithms are impractical if values of n and d are large. In dealing with large-scale datasets, two considerations play a crucial role for performing cost-effective KNN queries. First aspect is the computational cost. This means computing the similarity between two items must be performed fast. The second aspect is the storage cost. To accelerate the query execution, the data (or data signatures) should fit into the Random Access Memory (RAM), thus a preprocessing step is usually required to reduce the dimensionality of the input data.

To satisfy the above two requirements, there has been a growing interest in mapping (hashing) massive datasets to similarity preserving binary codes (binary vectors). The aim is to find a mapping such that similar items in the original dataset are mapped to close binary codes [14], [18], [29], [48], [50], [52]. Incorporating binary codes can dramatically reduce computational and storage costs, since binary codes are storage efficient and computing their respective distances (such as the Hamming distance) is a cheap computational operation. Moreover, in many applications, items of the database are naturally described in terms of binary codes. For example, in the *Bag of Words* (BoW) representation of the documents and images, the presence or absence of information, captured in terms of a binary variable, quantifies the items [7], [24].

Besides the computational and storage benefits, binary codes can be treated as indices (memory addresses) of a hash table. Relying on this property, to solve the KNN problem, one can first populate a hash table with binary codes, and then search the nearby buckets of the query until K items are retrieved. If binary codes are used as indices into a hash table, and the Hamming distance is used as the measure of similarity, then the algorithm that solves the exact KNN problem is straightforward. Starting with a Hamming radius equal to zero, $r = 0$, at each step, the algorithm checks all buckets that lie at the Hamming distance r from the query. After each step, r is increased by

one, and the algorithm proceeds until K items are retrieved. However, in many applications, binary codes are compared in terms of the cosine similarity, rather than the Hamming distance. This is known as the *angular KNN* problem. In such cases, there are no exact sequential procedures for checking the buckets of the hash table. In practice, instead of using a hash table, researchers resort either to the brute force search [14], or to approximate similarity search techniques, such as LSH [20], to solve the angular KNN problem.

In this paper, we propose a sequential algorithm for performing exact angular KNN search among binary codes. Our approach iteratively finds the sequence of hash table buckets to check until K neighbours are retrieved. We prove that, using the proposed procedure, the cosine similarity between the query and the sequence of generated buckets' indices will be monotonically decreasing. This means, the larger is the cosine similarity between a bucket index and the query, the sooner the index will appear in the sequence.

A problem with using hash tables is that the performance of search degrades quickly when the number of empty buckets is much larger than the occupied ones. In such scenarios, exhaustive search is a faster alternative. *Multi-Index Hashing* (MIH) [17], [30], [35], [36] is a powerful technique to address this issue. The MIH technique hinges on dividing long codes into disjoint shorter codes to reduce the number of empty buckets. Motivated by the MIH technique proposed in [35], we develop the *Angular Multi-Index Hashing* technique to realize similar advantages. Empirical evaluation of our approach shows orders of magnitude improvement in the search speed in conjunction with large-scale datasets.

Given a binary query and a hash table populated with binary codes, the research questions that we address are as follows:

- RQ1: In what order should the buckets be checked to find the K nearest neighbours?
- RQ2: How can the MIH technique be applied to angular preserving binary codes?
- RQ3: What is the effect of the MIH technique on the query time?

In a nutshell, we first establish a relationship between the cosine similarity and the Hamming distance. Relying on this relationship, a fast algorithm for finding the ordering of buckets to be checked is introduced. This allows modifying the multi-index hashing approach such that it can be applied to angular preserving binary codes. We empirically show that using the proposed approach to solve angular KNN queries can greatly reduce storage and computational costs. A handful of sublinear techniques are known that can provide approximate solutions to the angular KNN problem [1], [8], [44]. However, to the best of our knowledge, there is no general exact technique that performs better than brute force search to solve the angular KNN in conjunction with high dimensional binary codes.

2 Motivating Example

Consider the Google image search where the goal is to rank and find relevant images to a query image. One approach to solve this problem is to incorporate text-based analysis. For example, to find images of the Eiffel Tower, rather than searching through the visual content of photos, images that appear on the web pages containing the key word "Eiffel Tower" are chosen. This can often result in a poor outcome since no image analysis is performed. Indeed, to address this issue, Google uses the VisualRank algorithm [25]. The general idea of the VisualRank algorithm is to first find all the images that match the query by analyzing their corresponding tags, meta data, image names, and other related features, using the PageRank algorithm [37]. This results in forming a pool of candidate results. The next step is to compute the similarity between such candidates and the query. Since it is computationally impractical to measure distances to all actual images, hashing methods are incorporated to reduce the search time. Although this method can decrease the computational complexity, the search time can still be in the order of several minutes. A possible approach to tackle such a problem is to use a hash table of binary codes to reduce the search time. In this article, we address the problem of searching in a hash table of binary codes relying on the cosine similarity measure.

3 Related Works

This section reviews some of the popular and practical solutions to the nearest neighbour problem that relates to this study. In particular, we first review the early exact solutions to this problem. Then, we discuss the recent data-independent and data-dependent approximate algorithms that can trade scalability for accuracy.

A classical paradigm to reduce the computational cost of the nearest neighbour search relies on tree-based indexing structures which offer logarithmic query time, $O(\log n)$, on the average. Perhaps the most known example of such techniques is the *kd-tree* [3] with a worst-case search time of $O(dn^{1-1/d})$. Following [3], many other tree-based indexing algorithms are proposed, such as *R-tree*, *R*-tree*, *SS-tree*, *SR-tree*, and *X-tree* (see [41] for an overview). However, these methods suffer from a phenomenon known as the *curse of dimensionality*. This means, as the number of dimensions increases, these methods quickly reduce to the exhaustive search. Typically, the *kd-tree* and its variants are efficient for dimensions less than 20 [12], [49].

To overcome the apparent difficulty of devising algorithms that find the exact solution of the KNN problem, there has been an increasing interest to resort to *approximate* solutions. In the approximate formulation of KNN problem know as the *Approximate Nearest Neighbour* (ANN) problem, the algorithm is allowed to return an item whose distance from the query is at most c times ($c > 1$) the distance corresponding to the nearest neighbour. Among all such approximate techniques to solve the KNN problem, hashing is perhaps the most notable example. Hashing methods aim at mapping similar items into similar hash buckets. Early endeavors in hashing focused on using random functions to map the data. Among hashing techniques, the *Locality Sensitive Hashing* (LSH) [20], and its variants [9], [13], [42], are well-known. The idea of LSH is to use similarity preserving hash functions such that, with a high probability, points that are near to each other are mapped to the same hash bucket. The choice of the hash function is dependent on the measure of similarity in the original space. Over the years, several

family of hash functions are defined for different similarity measures. LSH constructs several hash tables, with the goal that the query point and its nearest neighbour hash to the same bucket in at least one of the hash tables. To achieve an acceptable retrieval accuracy, the number of hash tables, as well as the number of buckets in each hash table, should be relatively large [50], [51]. Experimental studies indicate that LSH-based techniques need over a hundred [13] and sometimes several hundreds hash tables [5]. Since the size of the hash table is proportional to the number of data objects, in some applications, using LSH results in excessive memory cost and long query time. To reduce these costs, different heuristics have been proposed over the years. *Entropy-based* LSH [38] and *Mutli-Prob* LSH [28] are two well-known techniques to reduce the number of hash tables. The idea is to form some points near the query (through random perturbations) and use them as additional queries. These techniques can reduce the number of hash tables at the cost of increasing the query time.

To find *compact* codes, instead of using random projections, recent studies have aimed at finding data dependent mappings to reduce the number of required bits. Salakhutdinov and Hinton [40] have used *Restricted Boltzmann Machine* with several hidden layers as an auto-encoder to learn the underlying binary codes. Weiss et al. [50] have proposed an eigenvector formulation which aims at finding short similarity preserving binary codes, such that bits are uncorrelated and balanced. *Binary Reconstructive Embedding* (BRE) [27] uses a loss function that penalizes the difference between the Euclidean distance in the input space and the Hamming distance in the binary space. Norouzi et al. [33] have used a hinge-like loss function to learn the codes. They derived an upper bound on the loss function, and optimized it instead of the actual objective function, to reduce the computational complexity of the underlying optimization problem. In a follow up work, Norouzi et al [34] have used a triple objective function to learn the binary codes for supervised problems. The corresponding loss function learns a projection such that items with the same label are at small Hamming distances to each other. More recently, Carreira-Perpinan and Raziperchikolaei [6] have used auto-encoder to learn binary codes in order to minimize the reconstruction error.

Current literature abounds with techniques for learning similarity preserving mappings [4], [14], [18], [29], [43], [48], [50], [52]. Optimizing parameters of hash functions is highly dependent on both the distance measure in the input space and that of the binary space. For unlabeled data, if the Euclidean distance is chosen as the measure of similarity, binary codes are usually compared in terms of the Hamming distance (which is equal to the squared Euclidean distance in the binary space) [15], [27]. However, in many vision and text-related applications, it is common to represent data as a *Bag of Words* (BoW) or vectors of frequencies with non-negative entries. In such applications, input vectors are compared in terms of the cosine similarity, and consequently, some techniques use the cosine similarity of binary codes as the similarity measure in the binary space [14], [45]. In the latter case, the objective is to find an *angular preserving* mapping that maps the high dimensional data to binary codes such that the cosine similarity of vectors in the input space is preserved in the binary space. Researches have proposed various metrics (other than the Hamming distance and the cosine similarity) to compare binary codes. A notable example is the *Spherical Hamming Distance* [18].

The idea of using hash tables to avoid exhaustive search in ANN techniques has been studied in recent years. Babenko et al. [2] have proposed the *inverted multi-index* technique to solve the ANN problem where data points are estimated with codewords of multiple codebooks. This technique creates two hash tables by decomposing data vectors into two disjoint subvectors and hashing each subvector in one of the tables. The query is similarly decomposed into two subvectors and the search is performed in each hash table to find the corresponding nearest neighbour. However, the memory overhead of this technique grows rapidly with the number of hash tables, and typically, the technique is efficient when only two hash tables are used. More recently, Matsui et al. [30] have used multiple hash tables to reduce the search time, whereby the distance between items are approximated by codewords of multiple codebooks. Iwamura et al. [21] have proposed a non-exhaustive search algorithm based on the branch and bound paradigm.

4 Definitions and Problem Statement

Table 1 shows the notations used throughout this paper. Some further notations will be defined in Section 5.

Given a dataset $\mathcal{X} = \{\mathbf{x}_i \in \mathbb{R}^d\}_{i=1}^n$, and a query $\mathbf{q} \in \mathbb{R}^d$, the aim of Nearest Neighbor (NN) problem, also called the 1 Nearest Neighbour (1NN) problem, is to find the item in $\mathcal{X}$ that is the closest to $\mathbf{q}$:

$$\text{NN}(\mathbf{q}) = \arg\min_{\mathbf{x}\in\mathcal{X}} dis(\mathbf{q}, \mathbf{x}) \tag{1}$$

where $dis(.,.)$ is the distance between two items. The K Nearest Neighbour problem (KNN) is the generalization of 1NN, aiming to find the K closest items to the query. An item $\mathbf{x} \in \mathcal{X}$ is in the result of KNN($\mathbf{q}$), if and only if it satisfies the following condition:

$$|\{\mathbf{t} \in \mathcal{X} \text{ such that } dis(\mathbf{t}, \mathbf{q}) < dis(\mathbf{x}, \mathbf{q})\}| < K \tag{2}$$

where $|.|$ denotes the cardinality of a set. When dis is the Euclidean distance, the problem is called the Euclidean KNN, and when dis is the inverse of the cosine similarity, it is called the angular KNN. The focus of this paper is to efficiently solve the angular KNN among a dataset of binary codes.

Another related search problem is the R-near neighbour problem (RNN).

Definition 1. (R-near neighbour Problem) Given a query $\mathbf{q}$, find all the R-near neighbours of $\mathbf{q}$ in the dataset, where an item $\mathbf{x}$ is an R-near neighbour of $\mathbf{q}$ if the distance between $\mathbf{x}$ and $\mathbf{q}$ is at most R.

Similarly, if the Euclidean distance is used as the similarity measure, we call the problem the Euclidean R-near neighbour.

RNN and KNN problems are closely related. One way to tackle the Euclidean KNN problem in conjunction with a dataset of binary codes is to solve multiple instances of the Euclidean RNN problem. First, a hash table is populated with binary codes in the dataset $\mathcal{B} = \{\mathbf{b}_i \in \{0,1\}^p\}_{i=1}^n$.

TABLE 1: Notations

Symbol	Explanation
K	Number of nearest neighbours to retrieve
$\mathbf{q}$	Binary query vector
$\|\|.\|\|_u$	ℓ_u norm
$\mathcal{C}(\mathbf{q}, e)$	Hamming ball centred at $\mathbf{q}$ with radius e
p	Length of the binary codes
$\mathcal{H}_{\mathbf{q},\mathbf{b}_i}$	Hamming distance tuple between $\mathbf{q}$ and $\mathbf{b}_i$
$sim(\mathbf{q}, \mathbf{b}_i)$	Cosine of the angle between $\mathbf{q}$ and $\mathbf{b}_i$
$r_{0\to1}^{\mathbf{q},\mathbf{b}_i}$	# of bits in which $\mathbf{q}$ is 0 and $\mathbf{b}_i$ is 1
$r_{1\to0}^{\mathbf{q},\mathbf{b}_i}$	# of bits in which $\mathbf{q}$ is 1 and $\mathbf{b}_i$ is 0
$\|\|\mathbf{x}-\mathbf{y}\|\|_H$	Hamming distance between vectors $\mathbf{x}$ and $\mathbf{y}$
$\mathcal{B}$	Dataset of binary codes

Then, starting from a Hamming radius equal to zero, $R = 0$, the procedure increases R and then solves the R-near problem by searching among the buckets at the Hamming distance R from the query. This procedure iterates until K items are retrieved. Nevertheless, if cosine similarity is used, the order of buckets to check will not be the same as the case of the Hamming distance. Unlike the Hamming distance, the angle between two binary codes is not a monotonically increasing function of their Euclidean distance. In other words, if binary codes $\mathbf{b}_1$ and $\mathbf{b}_2$ satisfy $||\mathbf{q}-\mathbf{b}_1||_H > ||\mathbf{q}-\mathbf{b}_2||_H$, it does not necessarily lead to $sim(\mathbf{q}, \mathbf{b}_1) < sim(\mathbf{q}, \mathbf{b}_2)$ where $sim(\mathbf{x}, \mathbf{y})$ is the cosine of the angle between binary codes $\mathbf{x}$, $\mathbf{y}$, and $||\mathbf{x}-\mathbf{y}||_H$ denotes the Hamming distance between $\mathbf{x}$, $\mathbf{y}$. Thus, the order of buckets to check for solving the angular KNN is different from the case of the Euclidean KNN. In the next section, we propose an algorithm which efficiently generates the ordering required to solve the angular KNN problem.

5 Fast Cosine Similarity Search

Consider a function $W : \mathbb{R}^d \to \{0,1\}^p$ that hashes the items in the dataset $\mathcal{X}$ while preserving the cosine similarity. That is, the cosine similarity between $\mathbf{x}_i$ and $\mathbf{x}_j \in \mathbb{R}^d$ is as close as possible to the cosine similarity between $W(\mathbf{x}_i)$ and $W(\mathbf{x}_j)$. Now suppose W is applied to all of the items in $\mathcal{X}$ to create the dataset of binary codes $\mathcal{B} = \{\mathbf{b}_i \in \{0,1\}^p\}_{i=1}^n$. The proposed technique uses the binary codes in $\mathcal{B}$ to populate a hash table where each binary code is treated as the direct index into a hash bucket. The problem that we aim to tackle is finding the K closest binary codes (in terms of the cosine similarity) to the query. Evidently, for a given query $\mathbf{q}$, the binary code that yields the highest cosine similarity is $\mathbf{q}$ itself. Therefore, the first bucket to check is the bucket with index $\mathbf{q}$. The next bucket to check is the one with the second largest cosine similarity to $\mathbf{q}$, and so on. In the rest of this section, we propose an algorithm for efficiently finding such a sequence of buckets.

The cosine similarity of two binary codes can be computed using:

$$sim(\mathbf{q}, \mathbf{b}_i) = \frac{\langle \mathbf{q}, \mathbf{b}_i \rangle}{||\mathbf{q}||_2 ||\mathbf{b}_i||_2}, \tag{3}$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product and $||\cdot||_u$ denotes the ℓ_u norm. In comparison to the Hamming distance, computing the cosine similarity is computationally more demanding. While computing the Hamming distance requires an XOR followed by the popcount operator, calculating the cosine similarity needs computing the square roots and a division, refer to (3). As mentioned, not all items with a fixed Hamming distance to the query have equal cosine similarities to the query. However, this article shows a remarkable connection between the Hamming distance and the cosine similarity of binary codes.

The set of all binary codes at the Hamming distance r from the query can be partitioned into $r+1$ subsets, where the codes in each subset yield equal cosine similarities to the query. In particular, for a binary code $\mathbf{b}_i$, with the Hamming distance r from $\mathbf{q}$, $\sum_{j=1}^{p} |\mathbf{q}^j - \mathbf{b}_i^j| = r$, the positions at which $\mathbf{q}$ and $\mathbf{b}_i$ differ can be one of the following two types: (1) $\mathbf{q}^j = 1$ and $\mathbf{b}_i^j = 0$, or (2) $\mathbf{q}^j = 0$ and $\mathbf{b}_i^j = 1$ where $\mathbf{x}^j$ denotes the j-th entry of the vector $\mathbf{x}$. Now, let $r_{1\to0}^{\mathbf{q},\mathbf{b}_i}$ and $r_{0\to1}^{\mathbf{q},\mathbf{b}_i}$ denote the number of type (1) and the number of type (2) distances, respectively. It is clear that $0 \le r_{1\to0}^{\mathbf{q},\mathbf{b}_i} \le ||\mathbf{q}||_1$ and $0 \le r_{0\to1}^{\mathbf{q},\mathbf{b}_i} \le p - ||\mathbf{q}||_1$. Also, the sum of type (1) and type (2) distances is equal to the Hamming distance, i.e., $r_{1\to0}^{\mathbf{q},\mathbf{b}_i} + r_{0\to1}^{\mathbf{q},\mathbf{b}_i} = r$. Consequently, we can write:

$$sim(\mathbf{q}, \mathbf{b}_i) = \frac{||\mathbf{q}||_1 - r_{1\to0}^{\mathbf{q},\mathbf{b}_i}}{\sqrt{||\mathbf{q}||_1} \times \sqrt{||\mathbf{q}||_1 - r_{1\to0}^{\mathbf{q},\mathbf{b}_i} + r_{0\to1}^{\mathbf{q},\mathbf{b}_i}}}. \tag{4}$$

The dot product of two binary codes (the numerator of (3)) is equal to the number of positions where $\mathbf{q}$ and $\mathbf{b}_i$ are both 1, which is equal to $||\mathbf{q}||_1 - r_{1\to0}^{\mathbf{q},\mathbf{b}_i}$. The denominator simply contains the ℓ_2 norms of $\mathbf{q}$ and $\mathbf{b}_i$. In the rest of this paper, we use (4) to compute the cosine similarity. For a given $\mathbf{q}$, each binary code $\mathbf{b} \in \mathcal{B}$ is associated with the tuple $(r_{1\to0}^{\mathbf{q},\mathbf{b}_i}, r_{0\to1}^{\mathbf{q},\mathbf{b}_i})$, called hereafter the *Hamming distance tuple*, which is defined as follows:

Definition 2. (Hamming Distance Tuple) Given a query $\mathbf{q}$, we say a given binary code $\mathbf{b}_i$ lies at the Hamming distance tuple $\mathcal{H}_{\mathbf{q},\mathbf{b}_i} = (r_{1\to0}^{\mathbf{q},\mathbf{b}_i}, r_{0\to1}^{\mathbf{q},\mathbf{b}_i})$ from $\mathbf{q}$ if:

1) number of bit positions in which $\mathbf{q}$ is 1 and $\mathbf{b}_i$ is 0 equals $r_{1\to0}^{\mathbf{q},\mathbf{b}_i}$, and,
2) number of bit positions in which $\mathbf{q}$ is 0 and $\mathbf{b}_i$ is 1 equals $r_{0\to1}^{\mathbf{q},\mathbf{b}_i}$.

A Hamming distance tuple, such as (r_1, r_2), is valid if both of its elements are in valid ranges, i.e., $0 \le r_1 \le ||\mathbf{q}||_1$ and $0 \le r_2 \le p - ||\mathbf{q}||_1$.

A key observation is that all binary codes with the same Hamming distance tuples yield the same sim value. Therefore, instead of searching for the sequence of buckets to be checked in the hash table, one can search for the sequence of Hamming distance tuples.

Each Hamming distance tuple represents a set of binary codes at the same angle from $\mathbf{q}$. The number of binary codes lying at the Hamming distance tuple $(r_{1\to0}^{\mathbf{q},\mathbf{b}_i}, r_{0\to1}^{\mathbf{q},\mathbf{b}_i})$ from $\mathbf{q}$ is:

$$\binom{||\mathbf{q}||_1}{r_{1\to0}^{\mathbf{q},\mathbf{b}_i}} \times \binom{p - ||\mathbf{q}||_1}{r_{0\to1}^{\mathbf{q},\mathbf{b}_i}}. \tag{5}$$

We say that a Hamming distance tuple (r_1', r_2') is less than or equal to (r_1, r_2), shown by $(r_1', r_2') \preceq (r_1, r_2)$, if and only if $r_1' \le r_1$ and $r_2' \le r_2$.

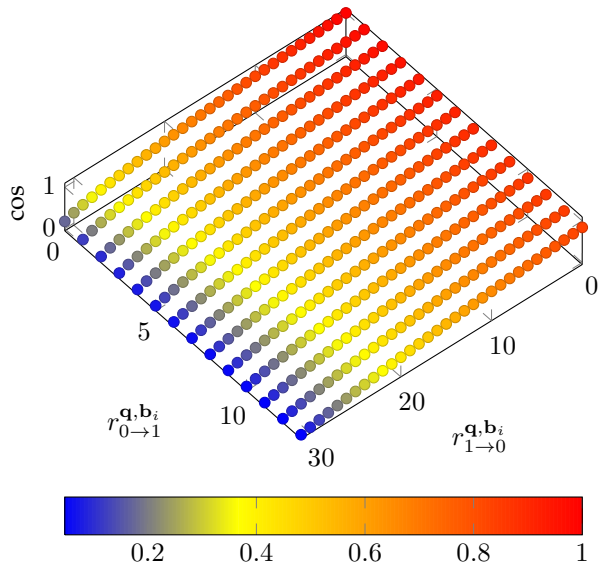

Fig. 1: Plot of sim values for different values of $r^{\mathbf{q},\mathbf{b}_i}_{1\to0}$ and $r^{\mathbf{q},\mathbf{b}_i}_{0\to1}$ with $p = 45$ and $||\mathbf{q}||_1 = 32$.

Definition 3. ((r_1, r_2))-NEAR NEIGHBOUR) A binary code $\mathbf{b}_i$ is called an (r_1, r_2)-near neighbour of $\mathbf{q}$, if we have $\mathcal{H}_{\mathbf{q},\mathbf{b}_i} \preceq (r_1, r_2)$.

Example 1. Suppose $\mathbf{q} = (1, 1, 1, 0, 0, 0)$ and $\mathbf{b}_1 = (0, 1, 0, 1, 1, 1)$, then $\mathbf{b}_1$ lies at the Hamming distance tuple $\mathcal{H}_{\mathbf{q},\mathbf{b}_1} = (2, 3)$ from $\mathbf{q}$, and $\mathbf{b}_2 = (1, 1, 1, 1, 1, 1)$ lies at the Hamming distance tuple $\mathcal{H}_{\mathbf{q},\mathbf{b}_2} = (0, 3)$ from $\mathbf{q}$. Also, the Hamming distance tuple $(0, 3)$ is less than the Hamming distance tuple $(2, 3)$.

The partial deviates of (4) with respect to $r^{\mathbf{q},\mathbf{b}_i}_{1\to0}$ and $r^{\mathbf{q},\mathbf{b}_i}_{0\to1}$ are both negative. This property indicates that, for a given Hamming distance tuple (x, y), all the binary codes with the Hamming distance tuple (x', y') satisfying $(x', y') \preceq (x, y)$ have larger sim values.

To visualize how the value of sim varies with respect to $r^{\mathbf{q},\mathbf{b}_i}_{1\to0}$ and $r^{\mathbf{q},\mathbf{b}_i}_{0\to1}$, the sim value as a function of $r^{\mathbf{q},\mathbf{b}_i}_{1\to0}$ and $r^{\mathbf{q},\mathbf{b}_i}_{0\to1}$ is plotted in Fig. 1. We are interested in sorting the tuples $(r^{\mathbf{q},\mathbf{b}_i}_{1\to0}, r^{\mathbf{q},\mathbf{b}_i}_{0\to1})$ (small circles in Fig. 1) in decreasing order of sim values. A straightforward way to do this, is to compute and sort the sim values of all possible tuples. However, we propose a more efficient approach that, in most cases, requires neither sorting, nor computing the sim values.

Definition 4. (Hamming Ball) For a given query $\mathbf{q}$, the set of all possible binary codes with a Hamming distance at most r from $\mathbf{q}$ is called the Hamming ball centered at $\mathbf{q}$ with radius r, and is shown by $\mathcal{C}(\mathbf{q}, r)$:

$$\mathcal{C}(\mathbf{q}, r) = \{\mathbf{h} \in \{0, 1\}^p : ||\mathbf{q} - \mathbf{h}||_H \leq r\}. \tag{6}$$

Given $\mathbf{b}_i$, values of $r^{\mathbf{q},\mathbf{b}_i}_{1\to0}$ and $r^{\mathbf{q},\mathbf{b}_i}_{0\to1}$ can be computed efficiently using bitwise operations. However, to search for the K closest neighbours in the populated hash table, we are interested in progressively finding the values of $r^{\mathbf{q},\mathbf{b}_i}_{1\to0}$ and $r^{\mathbf{q},\mathbf{b}_i}_{0\to1}$ that lead to binary codes with the largest sim value. One observation is that, within all indices lying at the Hamming distance r from $\mathbf{q}$, indices with the Hamming distance tuples $(r^{\mathbf{q},\mathbf{b}_i}_{1\to0}, r^{\mathbf{q},\mathbf{b}_i}_{0\to1}) = (0, r)$ and $(r^{\mathbf{q},\mathbf{b}_i}_{1\to0}, r^{\mathbf{q},\mathbf{b}_i}_{0\to1}) = (r, 0)$ (provided that they are valid tuples) yield the largest and the smallest cosine similarities with the query, respectively. This leads to the following proposition:

Proposition 1: Among all binary codes lying at the Hamming distance r from $\mathbf{q}$, those with larger values of $r^{\mathbf{q},\mathbf{b}_i}_{0\to1}$ yield larger sim values.

Proof: To prove this proposition, let us compute the derivatives of (4) with respect to $r^{\mathbf{q},\mathbf{b}_i}_{0\to1}$ or $r^{\mathbf{q},\mathbf{b}_i}_{1\to0}$. Suppose $r^{\mathbf{q},\mathbf{b}_i}_{1\to0} + r^{\mathbf{q},\mathbf{b}_i}_{0\to1} = r$, by replacing $r^{\mathbf{q},\mathbf{b}_i}_{1\to0}$ with $r - r^{\mathbf{q},\mathbf{b}_i}_{0\to1}$, we obtain:

$$sim(\mathbf{q}, \mathbf{b}_1) = \frac{||\mathbf{q}||_1 + r^{\mathbf{q},\mathbf{b}_i}_{0\to1} - r}{\sqrt{||\mathbf{q}||_1} \times \sqrt{||\mathbf{q}||_1 + 2r^{\mathbf{q},\mathbf{b}_i}_{0\to1} - r}}. \tag{7}$$

After some algebraic manipulations, it follows that $\frac{\partial sim}{\partial r^{\mathbf{q},\mathbf{b}_i}_{0\to1}} > 0$. Therefore, among all tuples at the Hamming distance r from $\mathbf{q}$, the maximum of sim occurs at $(r^{\mathbf{q},\mathbf{b}_i}_{1\to0}, r^{\mathbf{q},\mathbf{b}_i}_{0\to1}) = (0, r)$, and its minimum occurs at $(r^{\mathbf{q},\mathbf{b}_i}_{1\to0}, r^{\mathbf{q},\mathbf{b}_i}_{0\to1}) = (r, 0)$.■

Proposition 1 states that, for tuples at the Hamming distance r from $\mathbf{q}$, sim is a growing function of $r^{\mathbf{q},\mathbf{b}_i}_{0\to1}$. As a result, the order of tuples in the direction of decreasing sim values is $(0, r), (1, r-1), \ldots, (r, 0)$. In other words, among all the binary codes that lie at the Hamming distance r from the query, those with larger ℓ_1 norms yield larger cosine similarities.

While the Proposition 1 specifies the direction of the search for a given Hamming distance, it does not establish the relationship between the Hamming distance and the cosine similarity for different Hamming distances.

Example 2. For query $\mathbf{q} = (1, 0, 1, 0, 1)$, the binary code $\mathbf{b}_1 = (1, 0, 1, 0, 0)$ yields a smaller cosine similarity with $\mathbf{q}$ as compared to $\mathbf{b}_2 = (1, 1, 1, 1, 1)$ though it has a smaller Hamming distance to $\mathbf{q}$. As a result, the bucket corresponding to $\mathbf{b_2}$ should be checked before the one corresponding to $\mathbf{b}_1$.

Although the above may appear as a discouraging observation, we show that, for small Hamming distances, the cosine similarity and the Hamming distance are related to each other. In particular, the following proposition specifies the region where the cosine similarity is a monotonically decreasing function of the Hamming distance.

Proposition 2: If $||\mathbf{q}||_1 > \frac{r(r+t)}{t}$ for some $t \in \{1, \ldots, d\}$ and r, then all binary codes in $\mathcal{C}(\mathbf{q}, r)$ yield larger cosine similarities to $\mathbf{q}$ as compared to binary codes with Hamming distances at least $r + t$ from $\mathbf{q}$.

Proof: According to the Proposition 1, the maximum of the sim value for a fixed Hamming distance r occurs at $(r^{\mathbf{q},\mathbf{b}_i}_{1\to0}, r^{\mathbf{q},\mathbf{b}_i}_{0\to1}) = (0, r)$ with a sim value of $\sqrt{\frac{z}{z+r}}$, and its minimum occurs at $(r^{\mathbf{q},\mathbf{b}_i}_{1\to0}, r^{\mathbf{q},\mathbf{b}_i}_{0\to1}) = (r, 0)$ with a sim value of $\sqrt{\frac{z-r}{z}}$, where $z = ||\mathbf{q}||_1$. The condition in the proposition is satisfied if the smallest value of $sim(\mathbf{q}, \mathbf{b}_i)$, where $\mathbf{b}_i \in \mathcal{C}(\mathbf{q}, r)$, is larger than the largest value of $sim(\mathbf{q}, \mathbf{b}_j)$ where

$\mathbf{b}_j$ lies at the Hamming distance $r+t$ from $\mathbf{q}$. Hence, we have:

$$\begin{aligned}
&\sqrt{\frac{||\mathbf{q}||_1 - r}{||\mathbf{q}||_1}} > \sqrt{\frac{||\mathbf{q}||_1}{||\mathbf{q}||_1 + r + t}} \\
&\Rightarrow \frac{||\mathbf{q}||_1 - r}{||\mathbf{q}||_1} > \frac{||\mathbf{q}||_1}{||\mathbf{q}||_1 + r + t} \\
&\Rightarrow (||\mathbf{q}||_1 - r)(||\mathbf{q}||_1 + r + t) > ||\mathbf{q}||_1^2 \\
&\Rightarrow r(r+t) < t||\mathbf{q}||_1 \\
&\Rightarrow ||\mathbf{q}||_1 > \frac{r(r+t)}{t}.
\end{aligned} \tag{8}$$

This concludes the proof.■

If the condition of the Proposition 2 is satisfied for $t=1$ and an arbitrary value of r, then all the binary codes inside the Hamming ball $\mathcal{C}(\mathbf{q}, r)$ have larger cosine similarities as compared to those outside of $\mathcal{C}(\mathbf{q}, r)$. Also, among all binary codes inside $\mathcal{C}(\mathbf{q}, r)$, those with larger Hamming distances from the query have smaller cosine similarities. That is, if $\mathbf{b}_i$ is closer to $\mathbf{q}$ than to $\mathbf{b}_j$ in terms of the Hamming distance, then $\mathbf{b}_i$ is also closer to $\mathbf{q}$ in terms of the cosine similarity.

Therefore, for binary codes lying within the Hamming ball $\mathcal{C}(\mathbf{q}, r)$, cosine similarity is a decreasing function of the Hamming distance. In this case, the search algorithm is straightforward: for $t=1$, the maximum integer r that satisfies the inequality condition in the Proposition 2 is found. Let $\hat{r}$ denote the integer part of the positive root of the equation $r^2 + r - ||q||_1$ (this equation has only one positive root). Staring from $r=0$, the search algorithm increases the Hamming radius until the specified number of neighbours are retrieved, or until r reaches $\hat{r}$. For each Hamming radius, the search direction should be aligned with the direction specified by the Proposition 1.

DEFINITION 5. ((r_1, r_2)-NEAR NEIGHBOUR PROBLEM) Given the query $\mathbf{q}$ and the dataset of binary codes $\mathcal{B}$, the result of (r_1, r_2)-near neighbour problem is the set of all codes in $\mathcal{B}$ that lie at a Hamming distance tuple at most (r_1, r_2) from $\mathbf{q}$.

Our approach, outlined in Algorithm 1, is effective in cases when K binary codes are retrieved before r reaches $\hat{r}$. It tackles the angular KNN problem by solving multiple instances of the (r_1, r_2)-near neighbour problem. An important advantage of the proposed algorithm is that it does not need to compute the actual sim values between binary codes. It can be efficiently implemented using bitwise operators and the popcount function. In the rest of this section, we address the case of $r > \hat{r}$.

When the search radius is greater than $\hat{r}$, the sim value is not a monotonically decreasing function of the Hamming distance. However, we propose a sequential algorithm that can efficiently find the proper ordering of the tuples. Roughly speaking, at each step, the next tuple in the ordering can lie at many different Hamming distances. To tackle this issue, for each Hamming distance, among the tuples that are not selected so-far, we keep track of the tuple with the highest sim value. This procedure relies on a priority queue of Hamming distance tuples, where the priority of each tuple is determined by its corresponding sim value. The queue is initialized with the tuple $(0, \hat{r}+1)$. When a tuple is pushed into the queue, it is considered as traversed. At each subsequent step, the tuple with the top priority (the highest sim value) is popped from the queue. When a tuple is popped, two tuples are considered for insertion into the queue. Hereafter, these are called *the first anchor* and *the second anchor*, respectively. These two tuples are checked, and if "valid" and "not traversed", they are pushed into the queue.

Algorithm 1 Fast Cosine Similarity Search ($r \leq r'$)

Input: K: number of nearest neighbours to retrieve, $\mathbf{q}$: query, $\mathbf{H}$: hash table populated with binary codes
Output: $\mathcal{A}$: the set of retrieved items

```
Initialize set A = ∅
Initialize integer r = 0
Initialize r̂ with the positive root of the equation r^2 + r − ||q_1||
r̂ = ⌊r̂⌋
while |A| < K and r ≤ r̂ do
    R = (0, r)
    while R ≠ (r + 1, −1) do
        if R is a valid tuple then
            Check the buckets in H that lie at the Hamming distance tuple R from the query
            Add each of the found candidates to A
            R = R + (1, −1)
        end if
    end while
    r = r + 1
end while
```

DEFINITION 6. (First and Second Anchors of a Tuple) Given a query $\mathbf{q}$ and a Hamming distance tuple $R = (x, y)$, the first anchor and the second anchor of R are defined as follows:

- Among all tuples that lie at the Hamming distance $x+y+1$ from $\mathbf{q}$, the tuple with the largest sim value is called the *first anchor* of R.
- Among all tuples that lie at the Hamming distance $x+y$ from $\mathbf{q}$ and have smaller sim values as compared to R, the tuple with the largest sim value is called the *second anchor* of R.

EXAMPLE 3. For the query $\mathbf{q}$ with $||\mathbf{q}||_1 = 10$ and $p = 32$, the first anchor of $t = (1, 4)$ is $(0, 6)$ and the second anchor of $t = (1, 4)$ is $(2, 3)$ (according to the Proposition 1).

When a tuple is popped from the queue, the algorithm pushes the first and the second anchors of the popped tuple into the priority queue (provided that these are valid, and not traversed) and marks them as traversed. This procedure continues until either K elements are retrieved, or all valid tuples are traversed. Therefore, when a tuple such as $R = (x, y)$ is popped from the queue, the algorithm checks whether the following two tuples are valid or not:

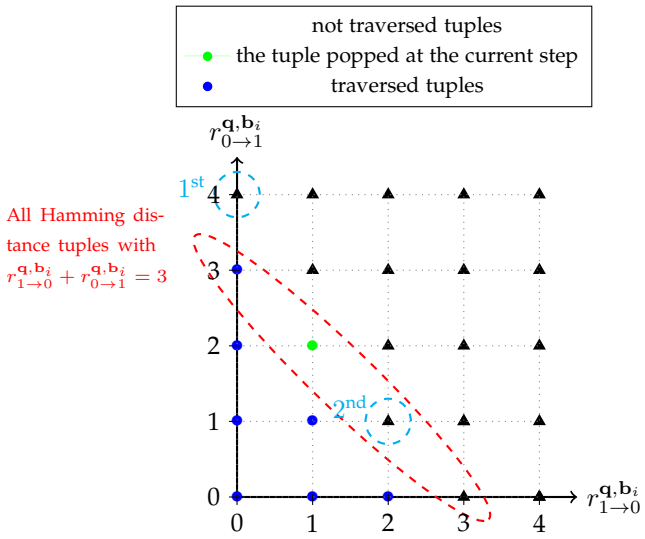

Fig. 2: Visual representation of the "first anchor" and the "second anchor" of a Hamming distance tuple.

1) **The first anchor of** R: This tuple, by definition, has the largest sim value among the tuples at the Hamming distance $x+y+1$ from the query. According to the Proposition 1, this candidate is $(0, x+y+1)$ if $x+y+1 \leq p - ||\mathbf{q}_1||$. In general, to ensure that the two components of this tuple are in acceptable ranges, the first anchor of R takes the form $(c, x+y+1-c)$ where $c = \max(0, x+y+1-(p-||\mathbf{q}||_1))$. Note that the first component of any Hamming distance tuple is at most $||\mathbf{q}||_1$ (number of ones in $\mathbf{q}$) and its second component is at most $p - ||\mathbf{q}||_1$ (number of zeros in $\mathbf{q}$).
2) **The second anchor of** R: Among the tuples that have smaller sim values as compared to R, and lie at the Hamming distance $x+y$ from the query, this tuple is the one that has the largest sim value. Using the Proposition 1, it is easy to show that the second anchor of R is $(x+1, y-1)$. This tuple is pushed into the queue if its components are in acceptable ranges (the second anchor is valid if $x+1 \leq ||\mathbf{q}_1||$ and $y-1 \geq 0$).

Figure 2 shows an example of the first/second anchors (shown in dashed circles) of a tuple that is selected in the current step (shown in green).

Algorithm 2 summarizes the steps of the proposed fast cosine similarity search in more details. Lines 9 to 22 are similar to Algorithm 1 and cover the case of $r \leq \hat{r}$. Lines 27 to 29 remove the tuple with the highest sim value from the priority queue and check its corresponding entries in the hash table. Lines 30 to 37 push the first and second anchors of the current tuple to the priority queue. The algorithm terminates when K binary codes are retrieved or when all tuples are checked.

Next, we prove that Algorithm 2 results in the correct ordering of Hamming distance tuples.

Proposition 3: In each iteration, the Hamming distance tuple popped from the queue has a smaller sim value as compared to the traversed tuples, and has the largest sim value among the not traversed tuples. Moreover, the algorithm eventually traverses every tuple.

Algorithm 2 Fast Cosine Similarity Search

Input: K: number of nearest neighbours to retrieve, $\mathbf{q}$: query, $\mathbf{H}$: hash table populated with binary codes
Output: $\mathcal{A}$: the set of retrieved items

1: Initialize set $\mathcal{A} = \varnothing$
2: Initialize integer $r = 0$
3: Create an empty priority queue pq
4: Marked all candidates as not traversed
5: Initialize $\hat{r}$ with the positive root of the equation $r^2 + r - ||\mathbf{q}||_1$
6: $\hat{r} = \lfloor \hat{r} \rfloor$
7: $R = (0, 0)$
8: **while** $|\mathcal{A}| < K$ **do**
9: **if** $r \leq \hat{r}$ **then**
10: **if** R is a valid tuple **then**
11: Check the buckets in $\mathbf{H}$ corresponding to the Hamming distance tuple R (i.e., perform a (r_1, r_2)-near neighbour)
12: Add the retrieved items to $\mathcal{A}$
13: Mark the tuple (r_1, r_2) as traversed
14: $R = R + (1, -1)$
15: **end if**
16: **if** $R = (r+1, -1)$ **then**
17: $r = r + 1$
18: $R = (0, r)$
19: **if** $r > \hat{r}$ **then**
20: $pq.push(R)$
21: **end if**
22: **end if**
23: **else**
24: **if** $pq.isempty()$ **then**
25: **return**
26: **end if**
27: $R \leftarrow pq.pop()$
28: Check the buckets in $\mathbf{H}$ corresponding to the Hamming distance tuple R
29: Add the retrieved items to $\mathcal{A}$
30: **if** The first anchor of R is not traversed **then**
31: $pq.push$(the first anchor of R)
32: Mark the first anchor of R as a traversed
33: **end if**
34: **if** The second anchor of R is a valid tuple **then**
35: $pq.push$(the second anchor of R)
36: Mark the second anchor of R as traversed
37: **end if**
38: **end if**
39: **end while**

Proof: When $r < \hat{r}$, then the algorithm is identical to the Algorithm 1, and according to the Propositions 1 and 2, the ordering is correct. For $r \geq \hat{r}$, we show that the selected candidate has the highest cosine similarity among the remaining tuples.

Assume that the Algorithm 2 is not correct. Let R be the first tuple that the algorithm selects incorrectly. This means another tuple, such as $R' = (x', y')$, yields the highest sim value and it has not been pushed into the priority queue.

The reason is that, if R' had been pushed into the queue, then R' would have been popped from the queue instead of R. Let $r' = x' + y'$, meaning that r' is the Hamming distance between $\mathbf{q}$ and any binary code lying at the Hamming distance tuple (x', y') from the $\mathbf{q}$. Consider all binary codes that lie at the Hamming distance r' from $\mathbf{q}$. If there exists a tuple with the second component greater than y' that has not been traversed so far, then a contradiction occurs (this means R' does not yield the largest sim value). This stems from the fact that, at a fixed Hamming distance, tuples with larger second components have larger sim values (Proposition 1). As a result, y' yields the largest possible value among the not-traversed tuples lying at the Hamming distance r' from $\mathbf{q}$. However, we show that, this tuple should have been pushed into the priority queue in previous steps. One of the following cases may occur:

1) Until the current step, no Hamming distance tuple at the Hamming distance r' from $\mathbf{q}$ has been selected: According to the Proposition 1, all tuples with the Hamming distance $r' - 1$ that are in the set $\mathcal{L} = \{(a, b) | (a, b)$ is a valid tuple and, $a + b = r' - 1, a \leq x', b \leq y\}$ have larger sim values than R', therefore, all of them must have been selected prior to R' in the sequence. However, the first time that a tuple from $\mathcal{L}$ was popped, R' was pushed into the priority queue. R' is in fact the first anchor of all the tuples in $\mathcal{L}$, and thus, it must have been pushed when any of the elements in $\mathcal{L}$ was popped from the priority queue.
2) At least one Hamming distance tuple with the Hamming distance r' from $\mathbf{q}$ has been traversed in previous steps: Similar to the previous case, R' was pushed into the priority queue when the algorithm popped the tuple $(x' - 1, y' + 1)$. In this scenario, R' is the second anchor of $(x' - 1, y' + 1)$.

It is concluded that R' must have been pushed into the priority queue during previous steps, which contradicts the assumption that R' is not a member of the priority queue.

We also need to prove that the algorithm is *complete*, i.e., the algorithm will not terminate until it either finds the K neighbours, or it traverses all the valid tuples. Again, let us assume the contrary. This means that, at the final step, the algorithm pops the last tuple from the queue and the last tuple does not have any valid anchors. Thus, the queue remains empty and the algorithm will terminate while there are some valid tuples that have not been traversed. It is clear that the not-traversed tuples cannot lie at the Hamming distance of r when at least one tuple with the Hamming distance r is traversed. The reason is that, once the first tuple with the Hamming distance r is popped from the queue, the second anchor of this tuple is pushed into the queue. Therefore, always one tuple with the Hamming distance r is in the queue until the last one of such tuples is popped, and such a last tuple does not have a valid second anchor. As a result, the only possible case is that all the tuples with a Hamming distance less than or equal to r have been traversed; and all of the tuples at a Hamming distance $r + 1$ and greater have not been traversed. However, this is not possible. The reason is that, when a tuple at the Hamming distance r is popped from the queue, its first anchor is pushed into queue and such a tuple should lie the Hamming distance $r+1$. Hence, all the tuples at the Hamming distance $r+1$ from the query will be covered eventually. This verifies the completeness of the algorithm.■

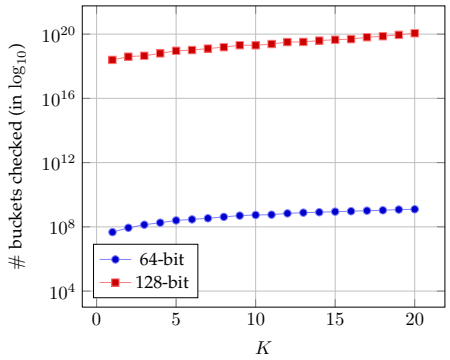

Fig. 3: The average number of required bucket checks required to solve the angular KNN problem, if a single hash table is used for the SIFT dataset (with 10^9 items).

6 Angular Multi-index hashing

Many applications require binary codes with large lengths (e.g., 64 bits) to achieve satisfactory retrieval accuracy. For such applications, it is not practical to use a single hash table with direct indices. For example, if 64-bit binary codes are used, then a hash table with 2^{64} buckets is needed, which typically cannot be fitted in the RAM.

A more crucial problem (in using larger binary codes with a single hash table) is the computational cost of the search. For long binary codes, we have $n \ll 2^p$ and many of the buckets will be empty. As a consequence, to solve the KNN problem, even for small values of K, often the number of buckets to be examined exceeds the number of items in the dataset. This means that the exhaustive search (linear scan) is a faster alternative as compared to using a hash table. As Fig. 3 shows, the average number of buckets that should be checked to solve the angular KNN query for the SIFT dataset, which contains one billion items (details of SIFT will be explained later), is often more than the number of available binary codes in the dataset. This problem arises as the number of buckets to be examined (many of which are empty) grows near-exponentially with the values of $r_{1\to0}^{\mathbf{q},\mathbf{b}_i}$ and $r_{0\to1}^{\mathbf{q},\mathbf{b}_i}$ (refer to (5)). In practice, to achieve a reasonable retrieval accuracy, the length of binary codes typically exceed 64 bits . To overcome this problem, recent studies have often resorted to linear search for binary codes longer than 64 bits [14], [40], [47].

Multi-Index Hashing (MIH) [17], and its variants [35], [36], are elegant approaches for reducing storage and computational costs of the R-near neighbour search for binary codes. The key idea behind the multi-index hashing is that, as many of the buckets are empty, one can merge the buckets over different dimensions of the Hamming space. To do this, instead of creating one huge hash table, MIH creates multiple smaller hash tables with larger buckets, where each bucket may be populated with more than one item. To do this, all binary codes are divided into smaller disjoint (usually with the same length) substrings, then each

substring is indexed within its corresponding hash table. For example, for $p = 64$, MIH may divide the binary codes into four 16-bit substrings, and then index each one into a hash table with 2^{16} buckets. Therefore, the number of buckets reduces to 4×2^{16}. More generally, if the number of hash tables is denoted by m, MIH reduces the number of required buckets from 2^p to $m2^{\frac{p}{m}}$ (assuming tables with direct addressing are used).

More importantly, MIH reduces the computational cost of the search. To solve the R-near neighbour problem, the query is similarly partitioned into m substrings. Then, MIH solves m instances of the $\frac{R}{m}$-near neighbour problem, one per each hash table. By doing this, the neighbours of each substring in its corresponding hash table are retrieved to form a set of potential neighbours. Since some of the retrieved neighbours may not be a true R-near neighbour, a final pruning algorithm is used to remove the false neighbours. Interested readers are referred to [36] for further details and cost analysis.

Despite of being efficient in storage and search costs, MIH cannot be applied to the angular preserving binary codes, since it is originally designed to solve the R-near neighbour problem in the Hamming space. In the rest of this section, we propose our *Angular Multi-index Hashing* (AMIH) technique for fast and exact search among angular preserving binary codes. Similar to many other search algorithms, AMIH includes two phases: (1) the indexing phase, and (2) the query phase. In the following, these two phases will be discussed in more details.

Instead of populating one hash table with binary codes, AMIH creates multiple smaller hash tables. To populate such smaller hash tables, each binary code $\mathbf{b} \in \{0,1\}^p$ is partitioned into m disjoint substrings $\mathbf{b}^{(1)}, \dots, \mathbf{b}^{(m)}$. For the sake of simplicity, in the following, we assume that p is divisible by m and use the notation $w = \frac{p}{m}$. As a result, the s-th hash table, $s \in \{1, \dots, w\}$, is populated with $\mathbf{b}_i^{(s)}$ $i \in \{1, \dots, n\}$. To retrieve the (r_1, r_2)-near neighbours of the query, $\mathbf{q}$ is similarly partitioned into m substrings, $\mathbf{q}^{(1)}, \dots, \mathbf{q}^{(m)}$.

The following proposition establishes the relationship between the Hamming distance tuple of two binary codes and their substrings.

Proposition 4: If $\mathbf{b}$ *lies at a Hamming distance tuple of at most* (r_1, r_2) *from* $\mathbf{q}$*, then:*

$$\begin{aligned} \exists \quad 0 < t \leq m \quad s.t. \quad & ||\mathbf{q}^{(t)} - \mathbf{b}^{(t)}||_H \leq \lfloor \frac{r_1 + r_2}{m} \rfloor \\ & \wedge \quad r_{1\to 0}^{\mathbf{q}^{(t)}, \mathbf{b}_i^{(t)}} \leq r_1 \\ & \wedge \quad r_{0\to 1}^{\mathbf{q}^{(t)}, \mathbf{b}_i^{(t)}} \leq r_2. \end{aligned} \tag{9}$$

The first condition follows from the Pigeon-hole principle. If in all of the m substrings, the Hamming distance is strictly greater than $\lfloor \frac{r_1+r_2}{m} \rfloor$, then we have $||\mathbf{q} - \mathbf{b}||_H \geq m(\lfloor \frac{r_1+r_2}{m} \rfloor + 1)$. This contradicts the assumption that $\mathbf{b}$ lies at a Hamming distance of at most $r_1 + r_2$ from $\mathbf{q}$. The second and the third conditions must in fact hold for all substrings. The reason is, if we have $r_{1\to 0}^{\mathbf{q}^{(t)}, \mathbf{b}_i^{(t)}} > r_1$, then we should have $r_{1\to 0}^{\mathbf{q}, \mathbf{b}_i} > r_1$. Similarly, if we have $r_{0\to 1}^{\mathbf{q}^{(t)}, \mathbf{b}_i^{(t)}} > r_2$, then we should have $r_{0\to 1}^{\mathbf{q}, \mathbf{b}_i} > r_2$. This means $\mathbf{b}$ is not a (r_1, r_2)-near neighbour of $\mathbf{q}$.■

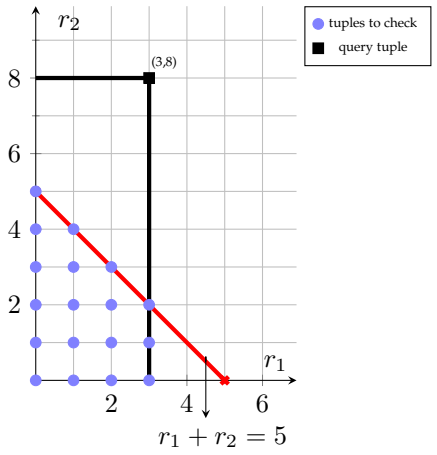

Fig. 4: The tuples that must be checked for solving the $(3, 8)$-near neighbour problem with 2 hash tables.

In simple terms, the Proposition 4 states that, if $\mathbf{b}$ is a (r_1, r_2)-near neighbour of $\mathbf{q}$, then at least in one of its substrings t, $\mathbf{b}^{(t)}$ must be a (r'_1, r'_2)-near neighbour of $\mathbf{q}^{(t)}$, where $r'_1 + r'_2 \leq \lfloor \frac{r_1+r_2}{m} \rfloor$, $r'_1 \leq r_1$ and $r'_2 \leq r_2$.

6.1 (r_1, r_2)-near Neighbour Search Using Multi-index Hashing

So far, we have established the necessary condition that facilitate the search among substrings. At the query phase, to solve a (r_1, r_2)-near neighbour search, AMIH first generates the tuples that satisfy the conditions of the Proposition 4. That is, to solve the (r_1, r_2)-near neighbour problem, AMIH generates the set of all tuples (r'_1, r'_2) such that $r'_1 + r'_2 \leq \lfloor \frac{r_1+r_2}{m} \rfloor$, where $r'_1 \leq r_1$ and $r'_2 \leq r_2$. This set is denoted by $\mathcal{T}_{r_1,r_2,m}$.

Example 4. Suppose $m = 2$ and we are interested in solving $(3, 8)$-near neighbour problem. According to the Proposition 4, we need to search among tuples with a Hamming distance at most $5 = \lfloor \frac{3+8}{2} \rfloor$ that satisfy the conditions in the Proposition 4. These tuples are shown in Fig. 4. Note that, for each tuple, one should check all its corresponding buckets in each of the hash tables.

Next, for each tuple such as $t = (r'_1, r'_2)$ in $\mathcal{T}_{r_1,r_2,m}$ and for each substring $\mathbf{q}^{(s)}$, $s \in \{1, \dots, m\}$, AMIH solves the (r'_1, r'_2)-near neighbour problem for the query $\mathbf{q}^{(s)}$ in the s-th hash table. This step results in a set of candidate binary codes, denoted by $\mathcal{O}_{j,t}$. According to Proposition 4, the set $\mathcal{O} = \bigcup_{j,t} \mathcal{O}_{j,t}$ is the superset of (r_1, r_2)-near neighbours of $\mathbf{q}$. Finally, AMIH computes the Hamming distance tuples between $\mathbf{q}$ and all candidates in $\mathcal{O}$, discarding the tuples that are not the true (r_1, r_2)-near neighbours of $\mathbf{q}$.

The intuition behind this approach is that, since the number of buckets that lie at the Hamming distance tuple (a, b) grows near-exponentially with the values of a and b, it is computationally advantageous to solve multiple instances of (a', b')-near neighbour problem with $a' < a$ and $b' < b$, instead of solving one instance of (a, b)-near neighbour problem where a and/or b are relatively large. This requires a significantly smaller number of buckets to be checked as compared to the case of a single large hash table. In what

follows, we show that AMIH can significantly reduce the query time in comparison with the linear scan.

6.2 Cost Analysis

The cost analysis directly follows the performance analysis of MIH in [36]. As suggested in [36], we assume that $\lfloor \frac{p}{\log_2 n} \rfloor \leq m \leq \lceil \frac{p}{\log_2 n} \rceil$. Using AMIH, the total cost per query consists of the number of buckets that should be checked to form the candidate set $\mathcal{O}$, plus the cost of computing the Hamming distance tuple between the retrieved binary codes in $\mathcal{O}$ and $\mathbf{q}$.

We start by providing an upper bound on the number of buckets that should be checked. Since the algorithm checks the identical buckets in each hash table, the number of bucket checks equals the product of m and the number of buckets that must be checked in a hash table.

To solve the (r_1, r_2)-near neighbour problem, for each tuple such as (a, b) in $\mathcal{T}_{r_1,r_2,m}$, the algorithm checks the buckets that correspond to (a, b). It is clear that, in the i-th hash table ($1 \leq i \leq m$), all binary codes corresponding to the tuples in the set $\mathcal{T}_{r_1,r_2,m}$ lie at a Hammming distance at most $\lfloor \frac{r_1+r_2}{m} \rfloor$ from the $\mathbf{q}^{(i)}$ (Proposition 4). Therefore, in the i-th hash table, the indices of buckets that must be checked are a subset $\mathcal{C}(\mathbf{q}^{(i)}, \lfloor \frac{r_1+r_2}{m} \rfloor)$, and we can write:

$$\begin{aligned} \#\text{buckets to check} &\leq \sum_{i=1}^{m} |\mathcal{C}(\mathbf{q}^{(i)}, \lfloor \frac{r_1+r_2}{m} \rfloor)| \\ &= \sum_{i=0}^{m} \sum_{j=0}^{\lfloor \frac{r_1+r_2}{m} \rfloor} \binom{p/m}{j} \\ &= m \times \sum_{j=0}^{\lfloor \frac{r_1+r_2}{m} \rfloor} \binom{w}{j} \\ &= m \times \sum_{j=0}^{\lfloor \frac{w(r_1+r_2)}{p} \rfloor} \binom{w}{j}. \end{aligned} \tag{10}$$

Assuming that $\frac{r_1+r_2}{p} \leq 1/2$, we can use the following bound on the sum of the binomial coefficients [11]:

For any $n \geq 1$ and $0 < \alpha \leq 1/2$, we have:

$$\sum_{i=0}^{\lfloor \alpha n \rfloor} \binom{n}{i} \leq 2^{H(\alpha)n}. \tag{11}$$

where $H(\alpha) := -\alpha \log(\alpha) - (1-\alpha)\log(1-\alpha)$ is the binary entropy of α.

Therefore, we can write:

$$\#\text{buckets to check} \leq m \sum_{j=0}^{\lfloor \frac{w(r_1+r_2)}{p} \rfloor} \binom{w}{j} \leq m 2^{wH(\frac{r_1+r_2}{p})}. \tag{12}$$

which is an upper bound on the number of buckets that must be checked. If binary codes are uniformly distributed in the Hamming space, the expected number of items per buckets of each hash table is $\frac{n}{2^w}$. Therefore, the expected number of items in the set $\mathcal{O}$ is:

$$E(|\mathcal{O}|) = \frac{n}{2^w} m \times 2^{wH(\frac{r_1+r_2}{p})}. \tag{13}$$

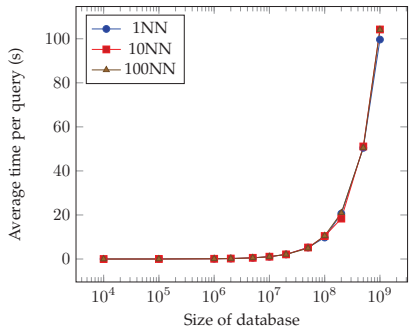

Fig. 5: Average linear scan search time per query on 64-bit binary codes.

Empirically, we observe that the cost of bucket lookup is only marginally higher than the cost of candidate test. If we have: *single lookup cost*= $t\times$ *single candidate test cost*, for some $t \geq 1$, then using (12) and (13), we can write the total cost as:

$$cost \leq m 2^{wH(\frac{r_1+r_2}{p})}(t + n/2^w). \tag{14}$$

For $m \approx p/\log_2 n$, by substituting $\log_2 n$ for w, we have:

$$cost = O(\frac{p}{\log_2 n} n^{H(\frac{r_1+r_2}{p})}). \tag{15}$$

For reasonably small values of $\frac{r_1+r_2}{p}$, the cost is sublinear in n. For example, for $\frac{r_1+r_2}{p} \leq 0.1$, the expected query cost would be $O(q\sqrt{n}/\log n)$.

The space complexity of AMIH comprises: 1) the cost of storing n binary codes each with p bits, which takes $O(np)$, and 2) the cost of storing n pointers to database items in each hash table. Each pointer can be represented in $O(\log_2 n)$ bits, therefore, the cost of storing pointers would be $O(mn \log_2 n)$. For $m = \lceil \frac{p}{\log_2 n} \rceil$, the total storage cost is $O(np + n \log_2 n)$.

7 Experiment and Results

We implemented the AMIH in C++ on top of the MIH implementation provided by the authors of [36] (all codes are compiled with GCC 4.8.4). The experiments have been executed on a 2.0 GHz Xeon CPU with 256 Gigabytes of RAM. Note that different results reported here are obtained using a single core to render the time cost comparisons meaningful. We first compare the performance of LSH with linear scan, and then provide an implementation-independent comparison between AMIH and LSH.

7.1 Datasets

In our experiments, we have used 2 non-synthetic datasets:
SIFT: The ANN_SIFT1B dataset [39] consists of SIFT descriptors. The available dataset is originally partitioned into 10^9 items as the base set, 10^4 items as the query set, and 10^8 items as the learning set. Each data item is a 128-dimensional SIFT vector.
TRC2: The TRC2 (*Thomas Reuters Text Research Collection 2*) consists of 1,800,370 news stories covering a period of one year. We have used 5×10^5 news as the learning set,

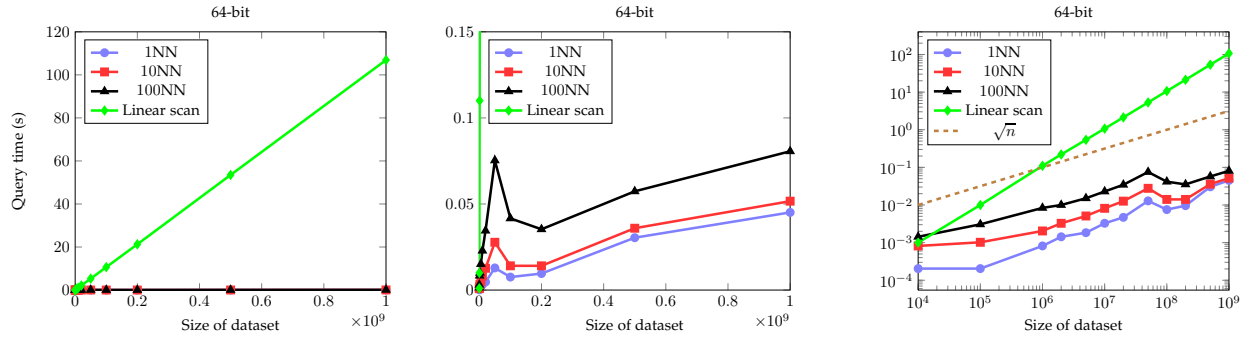

Fig. 6: Average search time for 64-bit binary codes of the SIFT dataset. AMIH and linear scan are executed to solve the KNN problem with $K \in \{1, 10, 100\}$.

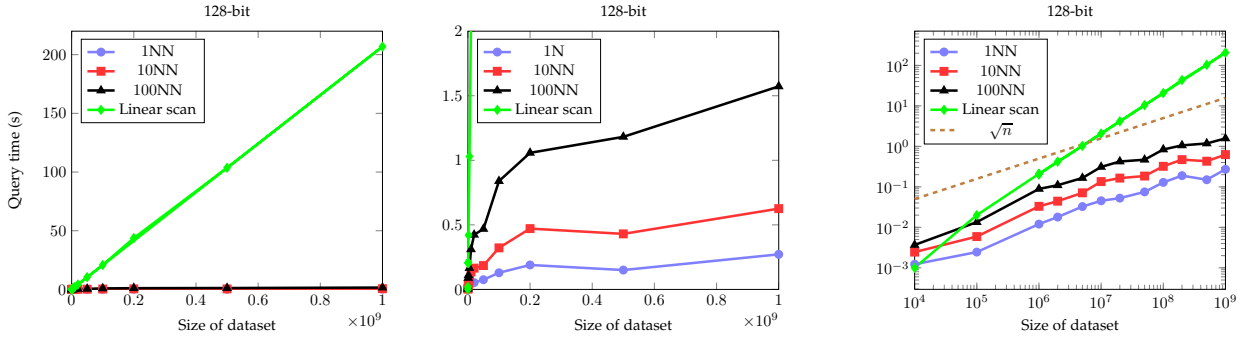

Fig. 7: Average search time for 128-bit binary codes of the SIFT dataset. AMIH and linear scan are executed to solve the KNN problem with $K \in \{1, 10, 100\}$.

10^6 news as the base set, and the remaining as the query set. We have preprocessed the data by removing common stop words, stemming, and then considering only the 2000 most frequent words in the learning set. Thus, each news is represented as a vector composed of 2000 word-counts.

We have used the angular-preserving mapping method called *Angular Quantization-based Binary Codes* (AQBC) proposed in [14], to create the dataset of binary codes. The MATLAB code for this method has been made publicly available by the authors of [14]. We have used the same code, with identical parameters and preprocessing steps, in all comparisons. For each dataset, the learning set is used to optimize the parameters of the hash function. Once learning is completed, the learning set is removed and the learned hash function is applied to the base and the query sets. The base set is used to populate the hash tables. Then, the angular KNN problem is solved for all queries in the query set and the average performance is reported.

7.2 AMIH and Linear Scan Comparison

The result of linear scan technique is compared with AMIH in terms of the search speed.

The linear scan algorithm for angular preserving binary codes searches among 10^6 binary codes in less than 1 second, and among 10^9 binary codes in almost 100 seconds (109 seconds, to be exact). The norm of any binary code with p bits ranges from 0 to $\sqrt{p}$. Thus, to increase the speed of the linear scan technique, we initialize a look up table with all the possible norm values. Moreover, as the term $\sqrt{||\mathbf{q}||_1}$ in the denominator of (4) is independent of $\mathbf{b}_i$, there is no need to account for its value in performing the search.

TABLE 2: Speedup gains that AMIH achieves in comparison to the linear scan.

		SIFT 1B		TRC2	
	# bits:	64	128	64	128
Speedup gain	1-NN	2672	1035	106	7.5
	10-NN	2137	345	27.5	3.21
	100-NN	1336	138	9.1	2.1
	Linear scan:	106	207	0.110	0.206

Figure 5 shows the average search speed per query using the linear scan technique for different database sizes. The plot indicates that the search time is almost independent of K (number of nearest neighbours). Consequently, for the sake of comparison, in the following, we use only the result of the linear scan technique for the 1NN problem. The linear scan technique can find the closest neighbour of a query within a dataset of 10^9 binary codes in almost 100 seconds, which is much faster as compared to using linear scan in the original space of SIFT vectors. Note that the linear scan technique can benefit from caching, as it performs sequential memory access. Otherwise, it would be much slower.

Figures 6 and 7 show the average query time as a function of the database size for 64-bits and 128-bits binary codes. The value of m (number of hash tables) for the

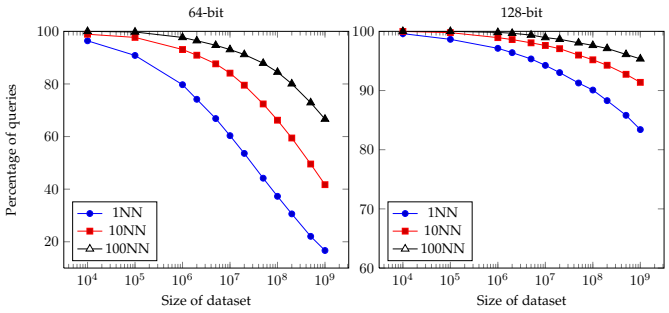

Fig. 8: The percentage of queries for which the Algorithm 2 enters the second case ($r > \hat{r}$).

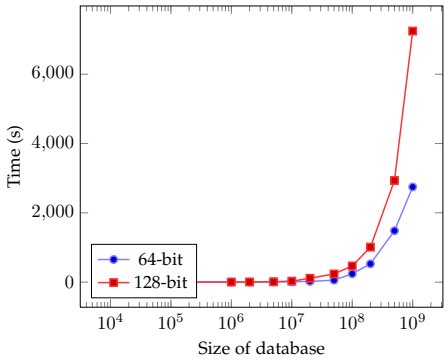

Fig. 9: Indexing time for AMIH on 10^9 SIFT dataset.

multi-index hashing technique in all experiments is set to $\frac{p}{\log_2 n}$, following [17], [30], [36]. The leftmost graphs show the search time in seconds in terms of the data base size. It is clear that the AMIH technique is significantly faster as compared to the linear scan technique for a broad range of dataset sizes and K values. To differentiate between the performance of the AMIH technique for different values of K, the middle graphs show the zoomed version of the leftmost graphs, and the rightmost graphs are plotted using logarithmic scale. As Figs. 6 and 7 illustrate, for the linear scan technique, the query time grows linearly with the database size, whereas the query time of the AMIH increases with the square root of the size. This means, the difference between the query times of the two techniques is more significant for larger datasets. For instance, the linear scan technique spends more than 3 minutes to report the nearest neighbour in the 10^9 SIFT dataset with 128-bits codes, while AMIH takes about a quarter of a second. The dashed line on log-log plots shows the growth rate of the $\sqrt{n}$ up to a constant factor. The evident similarity between the slope of this function and that of the AMIH query time indicates that, in practice, and even for non-uniform distributions, AMIH can achieve sublinear search time.

Figure 8 shows the percentage of queries for which Algorithm 2 enters the second case ($r > \hat{r}$). As the size of the database grows, the number of empty buckets reduces, and the algorithm finds the nearest neighbours within a smaller search radius. Similarly, for shorter binary codes, the number of buckets reduces, and in turn, Algorithm 2 can retrieve items before the search radius reaches $\hat{r}$.

Table 2 includes the speed up factors achieved by using the AMIH technique in comparison to the linear scan technique. Each entry in the table indicates the average query time using linear scan over the average query time using AMIH for a specific value of K. Interestingly, the AMIH technique solves the angular KNN problem up to hundreds and even thousands of times faster than the linear scan technique. In particular, AMIH can solve the 100NN problem 138 times faster than the linear scan on a dataset of 10^9 binary codes each with 128 bits. While the linear scan technique does not rely on any indexing phase, the AMIH technique requires each binary code to be indexed in m hash tables. The indexing time for AMIH using the SIFT dataset is shown in Fig. 9. For 64 and 128 bits codes, the indexing phase takes about 1 and 2 hours, respectively. However, this operation is performed off-line, and only once.

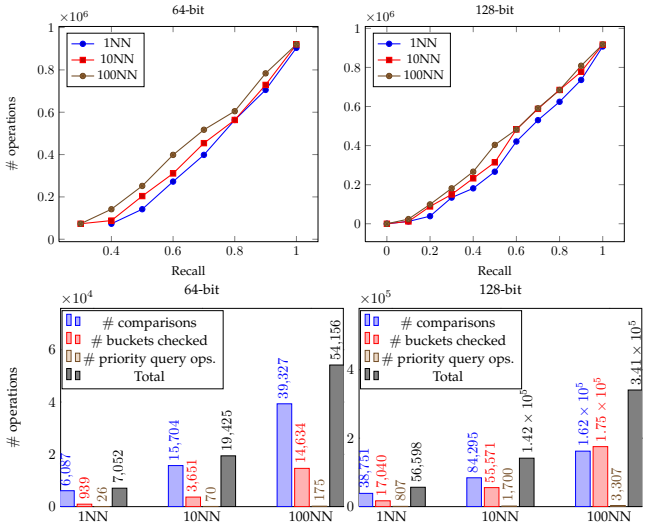

Fig. 10: Average number of operations required by LSH and AMIH to solve the KNN problem over 10^6 binary codes. The number of queries is 10^4.

7.3 AMIH and LSH Comparison

To the best of our knowledge, due to the curse of dimensionality, linear scan is the fastest exact technique for solving the general angular KNN problem. However, a handful of approximation algorithms exist that provide sublinear search time for this problem. The most well-known representative is the LSH which offers a (provably) sublinear search time. In the case of the angular KNN, LSH uses a family of hash functions that (with high probability) map close items (in terms of the a specific similarity measure) to the same bucket of a hash table. One of the popular hash functions for the cosine similarity measure is the *Simhash* function [8] (also known as the hyperplane LSH) which uses *sign random projection*. Given a vector $\mathbf{x}$, simhash utilizes a random vector $\mathbf{w}$ with each component generated from i.i.d normal distribution, $w_i \sim N(0,1)$, and only stores the sign of the projected data. Therefore, simhash is given by: $h_{\mathbf{w}}(\mathbf{x}) = sign(\langle \mathbf{w}, \mathbf{x} \rangle)$. Given such a hash function, LSH first hashes all the points in the base set, and then determines the near neighbours by hashing the query point and retrieving items stored in the buckets containing the query point. Finally, LSH performs a brute force search over the retrieved points to find the K closest items to the query.

Here, we compare the performance of LSH and AMIH for the task of angular KNN over binary codes. Similar to AMIH, LSH first checks the buckets of multiple hash tables and then performs a candidate check procedure. We compare LSH and AMIH in terms of the number of operations required to solve the angular KNN problem. This includes "checking the buckets of the hash tables" and "computing the distance between the retrieved items and the query". We also take into account the number of comparisons that the priority queue of AMIH requires. Comparing LSH and AMIH in this manner can provide a deeper insight regarding the performance of these techniques, as it is implementation-independent.

We have implemented the standard (G, L) parametrized LSH [20] algorithm with the simhash as the hash function. In particular, we have created L *compound* hash functions where a compound hash function such as c is formed by putting together G randomly selected hash functions, $c = [h_1, \ldots, h_G]$. In the preprocessing step, one hash table is created for each of the compound hash functions. At the query time, each query point is similarly hashed L times, and the elements stored in the buckets containing the query are retrieved. Then, the K closest points are selected by calculating the distance between the query point and all the retrieved items. To calculate the recall rate, the selected items by LSH algorithm are compared with the true K nearest neighbour of the query, and the percentage of true neighbours is reported. In the (G, L) LSH scheme, the ideal choice of G and L is dependent on the similarity thresholds and also on the hash function under consideration. An issue in using LSH is that the similarity threshold varies with the dataset. To ensure that the choice of these parameters does not effect our evaluations, we have tested all combinations of $G \in \{1, 2, \ldots, 20\}$ and $L = \{8, 16, 24, 32, 40, 48, 56\}$. We have used multiples of 8 for the number of bits in the compound hash functions. This is particularly attractive because the hash values fit into fixed number of bytes. For each combination of G and L, we have computed the mean recall of LSH for the task of KNN, together with the average number of points reported per query which is in fact the number of comparison operations that LSH must perform. We have then selected the least number of comparisons (over choices of K and L) required to achieve a given percentage of recall.

Figure 10 shows the average number of operations that AMIH and LSH perform. Both techniques are applied on 1 million binary codes; the original points are chosen randomly from the SIFT dataset, then AQBC algorithm [14] is applied to all points in the query and the base set. The top plots show the average recall rate in terms of the number of operations.

The bar charts indicate the number of operations required by the AMIH. In particular, three type of operations are shown: 1) comparisons between binary codes, 2) bucket check operations, and 3) the comparisons required by the priority queue–the *Total* bars indicate the sum of all operations. Note that, since AMIH is an exact search technique, the recall rate is always 1. According to Fig. 10, AMIH clearly exhibits orders of magnitude improvement in comparison to LSH. Even for relatively small values of the recall rate (say 70%), AMIH requires significantly fewer number of operations.

8 Conclusion and Future Work

This paper proposes a new algorithm for solving the angular KNN problem on large-scale datasets of binary codes. By treating binary codes as the memory addresses, our proposed algorithm can find similar binary codes in terms of the cosine similarity in a time that grows sublinearly with the size of dataset. To achieve this, we have first established a relationship between the Hamming distance and the cosine similarity. This relationship is in turn used to solve the angular KNN problem for applications where binary codes are used as the memory addresses of a hash table. However, using a hash table for long codes is often inferior to the linear scan due to the large number of empty buckets. To tackle this issue, as the second contribution, we have proposed the AMIH technique; a multi-indexing approach to reduce both computational and storage costs in comparison to using a single hash table. We have empirically showed that the AMIH technique can increase the search speed up to orders of magnitude when applied to large-scale datasets.

Most of our effort has focused on finding nearest neighbours of a binary code in a hash table when codes are compared in terms of the cosine similarity. There are a number of other similarity measures to compare binary codes. One potential avenue for future work is to find fast search algorithms for other measures of similarity such as the spherical Hamming distance [18] and the weighted Hamming distance [51].

While the search speed of the linear scan is only dependent on the size of the database, the AMIH search speed is dependent on the size, as well as the distribution of binary codes. Thus, applying AMIH to datasets with similar sizes would result in similar, but not identical, performances. Further, search speed analysis is required to reduce this threat. Also, the value of m in our experiments is chosen based on the value proposed for MIH in [36]. Since AMIH differs from MIH, this selection for the value of m may not be the optimal value for AMIH. Using cross validation methods to find better values for the parameter m can reduce this threat.

It has been shown that using asymmetric distance measures, such as the asymmetric Hamming distance, can boost the retrieval accuracy [16], [32]. Instead of hashing the query, these techniques compare the query point in the original space with the binary codes of the corresponding dataset. However, the higher accuracy often comes at the cost of a more complex search phase. One future line of work is to apply sublinear search algorithms for asymmetric hash functions.

References

[1] Alexandr Andoni, Piotr Indyk, Thijs Laarhoven, Ilya Razenshteyn, and Ludwig Schmidt. Practical and optimal lsh for angular distance. *Advances in Neural Information Processing Systems*, pages 1225–1233, 2015.

[2] Artem Babenko and Victor Lempitsky. The inverted multi-index. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3069–3076, 2012.

[3] Jon Louis Bentley. Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9):509–517, September 1975.

[4] Jonathan Brandt. Transform coding for fast approximate nearest neighbor search in high dimensions. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1815–1822, 2010.

[5] Jeremy Buhler. Efficient large-scale sequence comparison by locality-sensitive hashing. *Bioinformatics*, 17(5):419–428, 2001.

[6] Miguel A. Carreira-Perpinan and Ramin Raziperchikolaei. Hashing with binary autoencoders. *IEEE Conference on Computer Vision and Pattern Recognition*, 2015.

[7] Olivier Chapelle, Patrick Haffner, and Vladimir N Vapnik. Support vector machines for histogram-based image classification. *IEEE Transactions on Neural Networks*, 10(5):1055–1064, 1999.

[8] Moses S Charikar. Similarity estimation techniques from rounding algorithms. *ACM Symposium on Theory of Computing*, pages 380–388, 2002.

[9] Mayur Datar, Nicole Immorlica, Piotr Indyk, and Vahab S Mirrokni. Locality-sensitive hashing scheme based on p-stable distributions. *Annual Symposium on Computational Geometry*, pages 253–262, 2004.

[10] L Fei-Fei. Imagenet: crowdsourcing, benchmarking & other cool things. *CMU VASC Seminar*, 2010.

[11] J. Flum and M. Grohe. *Parameterized Complexity Theory*. Springer-Verlag Berlin Heidelberg, 2006.

[12] Jerome H. Friedman, Jon Louis Bentley, and Raphael Ari Finkel. An algorithm for finding best matches in logarithmic expected time. *ACM Transactions on Mathmatical Software*, 3(3):209–226, September 1977.

[13] Aristides Gionis, Piotr Indyk, Rajeev Motwani, et al. Similarity search in high dimensions via hashing. *International Conference on Very Large Data Bases*, 99:518–529, 1999.

[14] Yunchao Gong, Sanjiv Kumar, Vishal Verma, and Svetlana Lazebnik. Angular quantization-based binary codes for fast similarity search. *Advances in Neural Information Processing Systems*, pages 1196–1204, 2012.

[15] Yunchao Gong, Svetlana Lazebnik, Albert Gordo, and Florent Perronnin. Iterative quantization: A procrustean approach to learning binary codes for large-scale image retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(12):2916–2929, 2013.

[16] Albert Gordo, Florent Perronnin, Yunchao Gong, and Svetlana Lazebnik. Asymmetric distances for binary embeddings. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(1):33–47, 2014.

[17] Dan Greene, Michal Parnas, and Frances Yao. Multi-index hashing for information retrieval. *Annual Symposium on Foundations of Computer Science*, pages 722–731, 1994.

[18] Jae-Pil Heo, Youngwoon Lee, Junfeng He, Shih-Fu Chang, and Sung-Eui Yoon. Spherical hashing. *International Conference on Computer Vision and Pattern Recognition*, pages 2957–2964, 2012.

[19] Ling Hu, Wei-Shinn Ku, Spiridon Bakiras, and Cyrus Shahabi. Spatial query integrity with voronoi neighbors. *IEEE Transactions on Knowledge and Data Engineering*, 25(4):863–876, 2013.

[20] Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: Towards removing the curse of dimensionality. *Annual ACM Symposium on Theory of Computing*, pages 604–613, 1998.

[21] Mikio Iwamura, Takao Sato, and Kenji Kise. What is the most efficientway to select nearest neighbor candidates for fast approximate nearest neighbor search? *IEEE International Conference on Computer Vision*, pages 3535–3542, 2013.

[22] HV Jagadish, Johannes Gehrke, Alexandros Labrinidis, Yannis Papakonstantinou, Jignesh M Patel, Raghu Ramakrishnan, and Cyrus Shahabi. Big data and its technical challenges. *Communications of the ACM*, 57(7):86–94, 2014.

[23] H. Jegou, M. Douze, and C. Schmid. Product quantization for nearest neighbor search. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(1):117–128, 2011.

[24] Yu-Gang Jiang, Chong-Wah Ngo, and Jun Yang. Towards optimal bag-of-features for object categorization and semantic video retrieval. *International Conference on Image and Video Retrieval*, pages 494–501, 2007.

[25] Yushi Jing and Shumeet Baluja. Visualrank: Applying pagerank to large-scale image search. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(11):1877–1890, 2008.

[26] Mohammad R Kolahdouzan and Cyrus Shahabi. Continuous k-nearest neighbor queries in spatial network databases. *Spatio-Temporal Database Management*, pages 33–40, 2004.

[27] Brian Kulis and Trevor Darrell. Learning to hash with binary reconstructive embeddings. *Advances in Neural Information Processing Systems*, pages 1042–1050, 2009.

[28] Qin Lv, William Josephson, Zhe Wang, Moses Charikar, and Kai Li. Multi-probe lsh: efficient indexing for high-dimensional similarity search. *International Conference on Very Large Data Bases*, pages 950–961, 2007.

[29] Jonathan Masci, Michael M Bronstein, Alexander Bronstein, and Jurgen Schmidhuber. Multimodal similarity-preserving hashing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(4):824–830, 2014.

[30] Yusuke Matusi, Toshihiko Yamasaki, and Kiyoharu Aziawa. Pqtable: Fast exact asymmetric distance neighbor search for product quantization using hash tables. *International Conference on Computer Vision*, 2015.

[31] Jean-Baptiste Michel, Yuan Kui Shen, Aviva Presser Aiden, Adrian Veres, Matthew K Gray, Joseph P Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, Jon Orwant, et al. Quantitative analysis of culture using millions of digitized books. *Science*, 331(6014):176–182, 2011.

[32] Behnam Neyshabur, Nati Srebro, Ruslan R Salakhutdinov, Yury Makarychev, and Payman Yadollahpour. The power of asymmetry in binary hashing. *Advances in Neural Information Processing Systems*, pages 2823–2831, 2013.

[33] Mohammad Norouzi and David M Blei. Minimal loss hashing for compact binary codes. *International Conference on Machine Learning*, pages 353–360, 2011.

[34] Mohammad Norouzi, David M Blei, and Ruslan R Salakhutdinov. Hamming distance metric learning. *Advances in Neural Information Processing Systems*, pages 1061–1069, 2012.

[35] Mohammad Norouzi, Ali Punjani, and David J Fleet. Fast search in hamming space with multi-index hashing. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3108–3115, 2012.

[36] Mohammad Norouzi, Ali Punjani, and David J Fleet. Fast exact search in hamming space with multi-index hashing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(6):1107–1119, June 2014.

[37] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: Bringing order to the web. Technical Report 1999-66, Stanford InfoLab, November 1999.

[38] Rina Panigrahy. Entropy based nearest neighbor search in high dimensions. *Annual ACM-SIAM symposium on Discrete algorithm*, pages 1186–1195, 2006.

[39] Loïc Paulevé, Hervé Jégou, and Laurent Amsaleg. Locality sensitive hashing: A comparison of hash function types and querying mechanisms. *Pattern Recognition Letters*, 31(11):1348 – 1358, 2010.

[40] Ruslan Salakhutdinov and Geoffrey Hinton. Semantic hashing. *International Journal of Approximate Reasoning*, 50(7):969–978, 2009.

[41] Hanan Samet. *Foundations of multidimensional and metric data structures*. Morgan Kaufmann, 2006.

[42] Venu Satuluri and Srinivasan Parthasarathy. Bayesian locality sensitive hashing for fast similarity search. *International Conference on Very Large Data Bases*, 5(5):430–441, 2012.

[43] Jian Shao, Fei Wu, Chuanfei Ouyang, and Xiao Zhang. Sparse spectral hashing. *Pattern Recognition Letters*, 33(3):271–277, 2012.

[44] Anshumali Shrivastava and Ping Li. In defense of minhash over simhash. *International Conference on Artificial Intelligence and Statistics*, pages 886–894, 2014.

[45] Anshumali Shrivastava and Ping Li. Asymmetric minwise hashing for indexing binary inner products and set containment. *International Conference on World Wide Web*, pages 981–991, 2015.

[46] Antonio Torralba, Rob Fergus, and William T Freeman. 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(11):1958–1970, 2008.

[47] Antonio Torralba, Rob Fergus, and Yair Weiss. Small codes and large image databases for recognition. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008.

[48] Jun Wang, Sanjiv Kumar, and Shih-Fu Chang. Sequential projection learning for hashing with compact codes. *International Conference on Machine Learning*, pages 1127–1134, 2010.

[49] Roger Weber, Hans-Jörg Schek, and Stephen Blott. A quantitative analysis and performance study for similarity-search methods in high-dimensional spaces. *International Conference on Very Large Data bases*, 98:194–205, 1998.

[50] Yair Weiss, Antonio Torralba, and Rob Fergus. Spectral hashing. *Advances in Neural Information Processing Systems*, pages 1753–1760, 2009.

[51] Lei Zhang, Yongdong Zhang, Jinhu Tang, Ke Lu, and Qi Tian. Binary code ranking with weighted hamming distance. *IEEE*

Conference on Computer Vision and Pattern Recognition, pages 1586–1593, 2013.

[52] Lei Zhang, Yongdong Zhang, Jinhui Tang, Xiaoguang Gu, Jintao Li, and Qi Tian. Topology preserving hashing for similarity search. *International Conference on Multimedia*, pages 123–132, 2013.